Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

The development of creative search strategies

Yuval Hart^{a,*}, Eliza Kosoy^b, Emily G. Liquin^c, Julia A. Leonard^d, Allyson P. Mackey^d, Alison Gopnik^b

^a Psychology Department, Hebrew University of Jerusalem, Jerusalem, Israel

^b Psychology Department, University of California Berkeley, Berkeley, CA, USA

^c Psychology Department, Princeton University, Princeton, NJ, USA

^d Psychology Department, University of Pennsylvania, Philadelphia, PA, USA

ARTICLE INFO

Keywords: Creativity Children Exploration-exploitation Development Creative foraging game

ABSTRACT

How does creativity develop? Creativity is a multi-faceted behavior and thus it is difficult to find measures for creativity that are both precise and comparable across development. Here, we examine the development of creativity using a "creative foraging" task. The task measures different facets of creativity which we compare between 4- to 8-year-old children and adults. We find that compared to adults, children spend a higher percentage of their search exploring, and their exploitation phases are less efficient. Moreover, children orient their search to a different and smaller region of the search space, but within that space they produce more unique creative products. Lastly, as children grow up, their creative products become more adult-like and their uniqueness decreases. Together, these results suggest that creative search changes across development, in the search strategy employed, in how the space of possibilities is navigated, and in what ideas are ultimately chosen.

Creativity and exploration play a vital role in children's learning. Research has shown how sophisticated young children's exploration can be (Bonawitz, van Schijndel, Friel, & Schulz, 2012; Cook, Goodman, & Schulz, 2011; Legare, 2012; Schulz & Bonawitz, 2007; Schulz, Gopnik, & Glymour, 2007; Stahl & Feigenson, 2015; Ruggeri, Swaboda, Sim, & Gopnik, 2019) (see (Schulz, 2012) for a review). There is also some recent work suggesting that children may be more exploratory than adults in some cases (Blanco & Sloutsky, 2020; Liquin & Gopnik, 2020; Schulz, Wu, Ruggeri, & Meder, 2019; Sumner et al., 2019), corresponding to the common belief that children are more curious and more creative than adults. However, it is difficult to design tasks that are wideranging enough to capture genuine creativity-which we define as selfmotivated exploration through an unconstrained space of possibilities-and yet allow precise analyses and comparisons between children and adults. In the present research, we address this question by comparing both the dynamics and output of exploration in a creative search task between 4- to 8-year-old children and adults.

A literature across several disciplines and domains points to differences in possible exploration strategies and to an intrinsic tension between exploration and exploitation. In the optimality literature, for example, researchers have considered a wide range of problems that involve searching through a high-dimensional space for an optimal

solution (Kirkpatrick, Gelatt, & Vecchi, 1983). A "low-temperature" search may quickly settle on a locally optimal solution, furthering exploitation, but runs the risk of being stuck in a local minima. A more exploratory "high-temperature" search may help to reveal better solutions but also means searching through options that are less likely to be successful.

A related problem arises in the literature on Bayesian hypothesis search. Bayesian reasoning involves identifying the probability of different hypotheses given one's prior knowledge and the current evidence. However, for a reasonably complex problem, it is intractable to search through all the possible hypotheses, so instead various types of sampling methods can be used to select from the hypothesis space. Again, narrower sampling has the advantage of quickly arriving at a reasonable high probability hypothesis, whereas sampling more broadly allows more exploration of the space of potential hypotheses but means that the search process may take longer (Bonawitz, Denison, Griffiths, & Gopnik, 2014).

In reinforcement learning, an agent must balance exploration of the space of potential options or policies - sampling the space of possibilities to find the most rewarding options - with exploitation-taking advantage of those options that have been rewarding in the past (Cohen, McClure, & Yu, 2007; Daw, O'Doherty, Dayan, Seymour, & Dolan, 2006;

https://doi.org/10.1016/j.cognition.2022.105102

Received 22 May 2021; Received in revised form 11 February 2022; Accepted 18 March 2022

Available online 27 March 2022 0010-0277/© 2022 Published by Elsevier B.V.







^{*} Corresponding author. E-mail address: yuval.hart@mail.huji.ac.il (Y. Hart).

Gershman, 2018; Hills, Todd, Lazer, Redish, & Couzin, 2015; Teodorescu & Erev, 2014). In well-specified reinforcement learning problems such as multi-armed bandit problems, exploration is defined as the choices the agent makes to discover new beneficial actions that have not yet been tried, in contrast to the exploitation of known rewarding actions (Sutton & Barto, 2018).

Similarly, foraging studies suggest that foragers show two distinct phases - exploitation of resources that are clamped together, and once the patch of resources is depleted, going on exploration to find the next patch (Hills et al., 2015; Hills, Todd, & Goldstone, 2008; Pyke, 1984).

In all these cases, whether one is searching for the optimal solution, the most likely hypothesis, the most rewarding action or the most productive patch, an agent must choose whether to explore the space more widely or to narrow in more quickly on a reasonable solution. In many cases, the best strategy is to begin with wider exploration and then move to narrower exploitation.

Previous researchers (Gopnik, 2020; Gopnik, Griffiths, & Lucas, 2015; Lucas, Bridgers, Griffiths, & Gopnik, 2014) have suggested that young children explore the space of possibilities differently than adults do. In particular, children may come up with wilder ideas, may jump between ideas less systematically, or may explore the space of possibilities more widely. Providing preliminary support for this hypothesis, Lucas et al. (2014) presented 4- to 5-year-olds and adults with a Bayesian hypothesis inference problem. Participants saw evidence that a causal system operated on either a disjunctive principle (individual blocks caused a machine to activate) or a conjunctive principle (combinations of blocks caused a machine to activate). Following previous research (Lucas & Griffiths, 2010), adults learned the disjunctive relationship more successfully than the conjunctive relationship. However, children's learning was much more flexible. Like adults, children endorsed the disjunctive hypothesis after observing evidence supporting this rule, and like adults, they assumed the disjunctive hypothesis was correct when no evidence was provided. However, they were much more likely than adults to endorse the (a priori less likely) conjunctive hypothesis after they observed the corresponding evidence. This suggests that children's priors were broader and assumed more variability than adults, and that children accordingly sampled the hypothesis space in a more exploratory way (and see also (Gopnik et al., 2017; Wente et al., 2019)).

More recently, empirical work has begun to investigate developmental changes in explore-exploit decision making in the context of reinforcement learning (Christakou et al., 2013; Decker, Otto, Daw, & Hartley, 2016; Hauser, Iannaccone, Walitza, Brandeis, & Brem, 2015; Mata, Wilke, & Czienskowski, 2013; Nussenbaum & Hartley, 2019; Palminteri, Kilford, Coricelli, & Blakemore, 2016; Rovee & Rovee, 1969). In one study (Schulz et al., 2019), 7- to 11-year-olds performed more directed exploration (specifically sampling from uncertain regions in the search space (Auer, 2002; Frank, Doll, Oas-Terpstra, & Moreno, 2009; Wu, Schulz, Speekenbrink, Nelson, & Meder, 2018)) compared to adults, and they generalized more conservatively across the search space (see also Refs (Blanco & Sloutsky, 2020; Liquin & Gopnik, 2020; Nussenbaum & Hartley, 2019; Sumner et al., 2019)). In another study (Somerville et al., 2017), however, adolescents were less likely to forgo immediate reward for the sake of directed exploration compared to adults.

Although these studies begin to address developmental changes in exploration, they involve tasks in which exploration is pursued for an ultimate purpose in a constrained space, either discovering accurate causal hypotheses or maximizing reward. In contrast, creative discovery and real-world exploration are often spontaneous and intrinsically motivated (Amabile, 1985; Oh, Chesebrough, Erickson, Zhang, & Kounios, 2020), with a vast mental search space. Little is known about such self-motivated exploration in very large spaces of thought.

The general problem is an intrinsic tension between the breadth and spontaneity of wide-ranging creativity and the kind of precise analysis and control that is necessary to make comparisons across development. In the present research, we overcome this challenge using a "creative foraging" task, developed by Hart et al. (Hart et al., 2017; Hart et al., 2018), to precisely test the development of exploration in a more creative and spontaneous context. The task allows us to systematically investigate differences in the exploration patterns of younger and older children and adults on the same task. Do these differences involve the dynamics of the exploration itself, the search space over which children and adults choose to explore, the products of children's and adults' exploration, or some combination of all of these?

In the task, called the "Creative Foraging Game" (CFG), players use a tablet to create shapes from 10 connected squares, moving one square at a time, and choose shapes they deem "interesting and beautiful" (for adults) or "cool shapes or pictures" (for children) to a gallery (Fig. 1a). Creative search is thus intrinsically motivated, and participants choose when their search has resulted in an acceptable product (i.e. a shape deemed worthy of saving to the gallery). The game records players' moves, gallery choices, and the timing of their actions. The space of shapes in the game is large but well-defined and the distance between any two shapes (i.e., the shortest sequence of moves to get from one shape to another) is known.

The game measures not only the products of the creative search but also the dynamic patterns by which creative search proceeds. In previous research, Hart et al. (Hart et al., 2017) found that adults' search is composed of two phases – exploration and exploitation (see methods). In the exploration phase, players meander through the space of possibilities, moving from one shape to another on paths that are longer than the minimal path (Fig. 1b) – more like doodling than drawing a predetermined shape. On the other hand, during exploitation phases, players move on minimal paths, saving gallery shapes from a specific visual category (digits, letters, airplanes, etc., Fig. 1b), suggesting



Fig. 1. The creative foraging game measures players' search strategies and the regions of the search space they navigate in. a) Players create shapes from 10 connected squares by moving one square at a time. They can save to the "gallery" shapes they deem interesting and beautiful. The game records each player's individual moves, timings, and saved gallery shapes. b) For each two consecutive saved shapes we compute the efficiency score - dividing the number of moves of the minimal path between the two shapes (dark red arrows) by the number of moves of the player's actual path (grey arrows). We then calculate the average efficiency score in exploration and exploitation phases and the efficiency ratio between these two measures. c) We create the network of all exploitation clusters by connecting by an edge any two exploitation bouts that share at least two shapes. The resulting network has a giant component of connected communities that represent the common visual themes shared by players. For each player we calculate what percentage of their exploitation clusters belongs to these common themes. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

players visualize the next shape they wish to create before they act to create it. Prior research also tracked which kinds of shapes were produced, and how much shape categories overlapped across different participants. In particular, one can calculate what percentage of the categories in the exploitation phases belongs to the common themes shared by all participants (the giant component is found using the Girvan-Newman algorithm (Girvan & Newman, 2002), Fig. 1c). Among adults, there was a well-defined set of shape categories and those categories were very commonly shared across participants. Thus, the CFG task allows us to quantify the search strategies participants employ, the space they choose to navigate in, and the products of their search. Below, we compare the creative search performance of children and adults on this task and track the development of creative search through childhood.

As we describe above exploration may involve many different kinds of search in different domains. In the CFG paradigm we borrow from the foraging literature and consider exploration as the period of search for the next "patch" of related ideas. However, one key question addressed in this research is whether exploration in a space of creative ideas in fact resembles exploration as previously studied in Bayesian hypothesis learning, reinforcement learning, and foraging.

1. Methods

1.1. Participants

Participants were 146 children aged 4–8 years (mean \pm std. = 6 \pm 1.3 years). Data collection was conducted at the University of Pennsylvania, Princeton University, and University of California, Berkeley. The data from these child participants were compared to the adult participants (N = 100, Age: 20–49 years, mean \pm std. = 25 \pm 4 years) from Hart et al. (Hart et al., 2017). Twenty-eight additional child participants were excluded due to spending less than 8 min on the task or making less than 30 physical moves (a preregistered condition).

1.2. Pre-registration

The design, protocol, and analysis plan for this study were preregistered prior to data collection on the AsPredicted platform. We denote analysis that was not preregistered as exploratory in the Results section.

1.3. Procedure

Children played the Creative Foraging Game (Hart et al., 2017) on a tablet computer. An experimenter demonstrated to children how to move the squares on the screen (by tapping and dragging) and explained that all squares had to stay connected. Participants were instructed to save shape configurations that they thought made a "cool shape or picture" to the gallery. After these instructions, children spent a minimum of 8 min and a maximum of 12 min completing the task. The experimenter provided assistance with moving squares and saving shapes as needed, but all actions were child-directed. We captured the dynamics of exploration by recording each participant's saved shapes, as well as the sequences of actions taken (squares moved) to get from one saved shape to the next. The adult participants completed the task as described by Hart et al. (Hart et al., 2017).

1.4. Game measures analysis

Though we provide further detail in the Results, we briefly describe the main procedures used to analyze the creative search data. To analyze the search process, we used a thresholding algorithm to tag exploitation and exploration phases, defining exploitation phases as phases where players' time intervals between consecutive choices to the gallery are (weakly) monotonically decreasing, and exploration phases as periods where the time intervals between consecutive choices to the gallery are monotonically increasing. We calculate the efficiency of each exploration (exploitation) bout by dividing the minimal number of moves required between each two consecutive choices of shapes in exploration (exploitation) with the actual number of moves the player made. From these measures we can calculate players' efficiency ratio – the ratio of their averaged exploitation efficiency and averaged exploration efficiency. This efficiency ratio thus gauges individuals' general efficiency in exploitation as opposed to exploration.

To analyze the products of creative search, we build a network of "shared meaning" shapes: Each exploitation bout has a cluster of chosen shapes. We collect all the clusters of shapes from all players and connect clusters of shapes that share at least two shapes. On this network, we use the Girvan-Newman (Girvan & Newman, 2002) (community finder) algorithm to extract the common themes shared by players. Then, for each player we compute the percent of their exploitation phases that belong to these common themes (see Hart et al. (Hart et al., 2017) for more details). In summary, this scores the extent to which each participant produced shapes that align with the shapes commonly produced by many other participants.

2. Results

2.1. Children differ in their exploration and exploitation search behavior

We find that children (N = 146) alternated between phases of exploration and exploitation, though they took fewer total moves and saved fewer shapes than adults (p's $< 10^{-4}$). During their search, children spent a greater proportion of their time in the exploration phases relative to adults (Children: mean \pm ste = 0.23 \pm 0.01, 95% CI = [0.21, 0.27], Adults: mean \pm ste = 0.19 \pm 0.01, 95% CI = [0.16, 0.21]. Mann-Whitney test: U(14,746) = 6022, p < 0.014, Rank biserial correlation = 0.183, see SI, Fig. S1). Children also saved a greater proportion of their shapes in the exploration phases relative to adults (Children: mean \pm ste = 0.26 \pm 0.01, 95% CI = [0.24, 0.29], Adults: mean \pm ste =0.19 \pm 0.01, 95% CI = [0.17, 0.21]; Mann-Whitney test: *U*(14,746) = 5283.5, p < 0.0002, Rank biserial correlation = 0.283, Fig. 2a and Fig. 2a, Inset for the cumulative density function (CDF) differences between the two distributions). A shift-function (Rousselet, Foxe, & Bolam, 2016; Rousselet, Pernet, & Wilcox, 2017) analysis of the decile differences between the two distributions shows similar results (see SI). Thus, children both tend to explore more in the game and save a higher percentage of their shapes during their exploration phases.

Children also demonstrated their exploratory nature in their less efficient search during exploitation. For each two consecutive saved shapes, we calculated the ratio between the minimal number of moves required between the two saved shapes and the number of actual moves the player takes between these two shapes in her game (Fig. 1b). This efficiency score measures the efficiency of player's moves during exploration and exploitation. We find that during exploitation phases, children move between saved shapes in a less efficient way than adults do, as expressed in their exploitation efficiency scores (Children: mean \pm ste = 0.52 \pm 0.01, 95% CI = [0.49, 0.55], Adults: mean \pm ste = 0.57 \pm 0.02, 95% CI = [0.54, 0.60], Mann-Whitney U(14,746) = 8430.5, p < 10000.033). Although children's exploitation efficiency was smaller than adults, their exploration efficiency was comparable to adults (Children: mean \pm ste = 0.50 \pm 0.02, 95% CI = [0.47, 0.53], Adults: mean \pm ste = 0.47 ± 0.02 , 95% CI = [0.43,0.52], p > 0.209). To capture the differences between children and adults in both phases, we calculated the efficiency ratio, which computes the ratio between exploration efficiency and exploitation efficiency. We find that children's efficiency ratio (exploration efficiency/exploitation efficiency) was higher than adults, providing further support for the idea that children were less efficient in their exploitation phases (Children: mean \pm ste = 1.07 \pm 0.06, 95% CI = [0.97, 1.19], Adults: mean \pm ste =0.84 \pm 0.04, 95% CI = [0.77, 0.91]. Mann-Whitney test: U(14,746) = 5611, p < 0.003, Rank biserial correlation = 0.223, Fig. 2b). Together, these findings suggest

Y. Hart et al.



Fig. 2. Children are more exploratory than adults. a) Children save a higher proportion of gallery shapes during the exploration phase (Children: mean \pm ste = 0.26 \pm 0.01, 95% CI = [0.24, 0.29], Adults: mean \pm ste $=0.19 \pm 0.01$, 95% CI = [0.17, 0.21]. Mann-Whitney test: U (14,746) = 5283.5, p < 0.0002, Rank biserial correlation = 0.283). Inset, the differences in the CDF of the children's vs. adults' distributions of % shapes saved to the gallery in exploration. Error bars are 95% CI of each quantile difference, b) Children show higher efficiency ratios, defined as the ratio of their exploration efficiency vs. their exploitation efficiency, suggesting children exploitation is less effective than adults (Children: mean \pm ste = 1.07 \pm 0.06, 95% CI = [0.97, 1.19], Adults: mean \pm ste

 $=0.84 \pm 0.04$, 95% CI = [0.77, 0.91]. Mann-Whitney test: U(14,746) = 5611, p < 0.003, Rank biserial correlation = 0.223). Inset, the differences in the CDF of the children's vs adults' distributions of efficiency ratio values. Error bars are 95% CI of each quantile difference,

that children's search of the space of possibilities is different in character than that of adults - children both spend more time exploring, and they exploit less effectively. Given the different search strategies children employ, we next asked whether the space in which they navigate, and the products of their creative search are different as well.

2.2. Children navigate in a different and smaller part of the space of possible shapes

Since the space of possible shapes is well defined and metric, we can chart the different parts of the space children and adults explored. Interestingly, we find that children navigate in a different and smaller part of the space of possible shapes (See Fig. 3a for a depiction of the network of gallery shapes explored by children and adults and Fig. 3b for the top 15 most saved shapes for children and adults.) As a measure of children's and adults' different paths in the space of shapes, in an exploratory analysis we calculated the log-likelihood of each path according to adults' choices. Specifically, we computed the probability of each transition from one shape to another by counting the number of

moves between the two shapes across all adult participants, then dividing by the number of adults who made the first shape. We then calculated the mean of minus log-likelihoods of the probabilities along the entire path of each child and adult participant. Therefore, a higher score means lower probability of that transition in the adult participants. Children's paths in the space of possible shapes were different from those of adults, as shown in their log-likelihood scores (Children: mean \pm ste = 2.18 \pm 0.02, 95% CI = [2.14, 2.22], Adults: mean \pm ste = 1.42 \pm 0.02, 95% CI = [1.38, 1.46]. Mann-Whitney test: U(14,746) = 224, p $< 10^{-5}$, Rank biserial correlation = 0.97), Fig. 3c). To further support the idea that children may be navigating a space of possibilities that is distinct from the space of possibilities where adults focus their attention, and that that space is smaller in size, in an exploratory analysis we computed the percentage of shared shapes between children and adults out of the total number of shapes created by adults. To account for the lower number of moves children do in their search, we compared this percentage to a truncated version of adults' games, where the number of steps in adults' games was truncated to match the mean number of steps in children's games. We repeated this calculation 1000 times with







Fig. 3. Children's network of shapes is different and smaller than adults' network of shapes, a) The network of gallery shapes of adults and children. Each node is a shape saved to the gallery; each edge connects two consecutive choices of gallery shapes. Nodes colors represent - shapes chosen by both adults and children (Blue), shapes chosen only by adults (Green), and shapes chosen only by children (Orange). b) The top 15 most saved shapes from children's and adults' network of saved shapes. Shapes are grouped by adults (green), children (orange), and shared shapes (at the intersection of the two ellipses). c) Children's paths in the search space are different from adults' paths (Children: mean \pm ste = 2.18 \pm 0.02, 95% CI = [2.14, 2.22], Adults: mean \pm ste = 1.42 \pm 0.02, 95% CI = [1.38, 1.46]. Mann-Whitney test: $U(14,746) = 224, p < 10^{-5}$, Rank biserial correlation = 0.97). Inset, the differences in the CDF of the children's vs adults' distributions of paths' mean minus loglikelihood values. Error bars are 95% CI of each quantile difference, (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

а

sampling with replacements to bootstrap the distributions of the possible percentage of shared saved shapes. Children's percentage of shared shapes was significantly lower than that of the adults' truncated version (Children: mean \pm ste = 20% \pm 0.02%, 95% CI = [19%, 22%], Adults: mean \pm ste = 40% \pm 0.06%, 95% CI = [36%, 44%]. Mann-Whitney test: U(10⁶) =0, $p < 10^{-5}$, Rank biserial correlation = 1, see SI where we also show that removing Hebrew letters from adults' shapes does not change the results). These results support the hypothesis that children navigate in a smaller and distinct space of possibilities compared to adults.

2.3. Children show higher uniqueness in their creative search products than adults

We next compared the products of children's and adults' search: the shapes that participants chose to save to the gallery. We find that the nature of children's saved shapes was different than that of adults. Children saved a significantly higher fraction of unique shapes (shapes that only a single participant saved to the gallery) in their creative search than adults did (Children: mean \pm ste = 0.40 \pm 0.02, 95% CI = [0.37, 0.44], Adults: mean \pm ste = 0.27 \pm 0.02, 95% CI = [0.24, 0.31]. Mann-Whitney test: U(14,746) = 4737, $p < 10^{-5}$, Rank biserial correlation = 0.358, Fig. 4a. Similar results are obtained even when we truncate adults' trajectories to match children's average trajectory length, see SI). As a second measure of their individuality, we can calculate for each child (adult) the fraction of exploitation phases that belong to the common themes in children's (adults') search (see Methods, and Fig. 4b). We find that the fraction of categories children chose to exploit from their common themes was lower compared to adults (Children: mean \pm ste = 0.17 \pm 0.01, 95% CI = [0.14,0.20], Adults: mean \pm ste = 0.24 \pm 0.01, 95% CI = [0.21,0.27]. Mann-Whitney test: U(14,746) = 9307, p < 0.0004, Rank biserial correlation = -0.262, Fig. 4b).

2.4. As children develop, their creative search becomes closer to adults' search space

Our results suggest that children's creative search is different in its dynamic features, in the space of possibilities they navigate, and in their choice of products compared with the creative search of adults. Are there similar differences between younger and older children?

To study this question, we compared children's performance throughout the age range of 4–8 years. We separated children into two groups using a median split of age: 1) Younger children, younger than or equal to 6 years old and 2) Older children, older than 6 years old. Analysis was also made with age as a continuous variable, see SI.

Contrary to the differences between children and adults in their exploration-exploitation relation, we did not find significant changes in the dynamic features of children's creative search across this age range (neither in their proportion of time in exploration, nor in their efficiency ratios, all p's > 0.24). We did, however, find that as children grow, they save a smaller fraction of unique shapes in their creative search (Younger children: mean \pm ste $= 0.44 \pm 0.02$, 95% CI = [0.40, 0.49], Older children: mean \pm ste $= 0.35 \pm 0.03$, 95% CI = [0.30, 0.40]. Mann-Whitney test: U(5265) = 3287, p = 0.010, Rank biserial correlation = -0.249, Fig. 5a. Similar results are obtained when we correlate the fraction of unique shapes with children's absolute age in years, see SI).

In addition, in an exploratory analysis we measured how many saved shapes younger and older children shared with adults, indicating how children's search space compared with the adults' space. Accordingly, we calculated the percentage of saved shapes that were common to adults and children in each age group out of the number of shapes each group of children created. We repeated this calculation 1000 times with sampling with replacements to bootstrap the distributions of the possible percentage of shared saved shapes. Older children saved a higher percentage of shapes that were shared with adults than younger children (Younger children: mean \pm ste = 38.4% \pm 0.1%, 95% CI = [35.8%, 41.8%], Older children: mean \pm ste = 48.2% \pm 0.1%, 95% CI = [48.0%, 52.8%]. Mann-Whitney test: U(10⁶) = 10⁶, $p < 10^{-5}$, Rank biserial correlation = 0.99, Fig. 5b).

Combined, these findings suggest that as children grow older, their search space comes to resemble the adult search space. Children in the 4–8 year old range we studied do not yet show the dynamic features of adults' creative search (their exploration-exploitation balance and efficiency ratio), which might suggest that these dynamic features develop later in childhood.

2.5. Children's creative search behavior is different from random exploration

One possible explanation of children's exploratory behavior is simply that their search is random, and thus they spend more time in exploration, are less efficient in their exploitation, and have a larger fraction of unique shapes. To test whether this was true, in an exploratory analysis, we compared the distribution of common shapes in children's, and adults' behavior, to random exploration behavior. We can compare these distributions by noting that different exploration mechanisms will produce different scaling laws of the distribution of created shapes - A heavy tailed distribution (a small power-law) indicates an exploration process that is directed towards common shapes, while a fast-decaying distribution (a large power-law) indicates an exploration process that is not directed towards specific shapes.

We simulated random exploration by creating trajectories of random moves in the space of shapes with the total length of trajectories sampled from the distribution of lengths of children's games. Next, to simulate choices to the gallery in the random trajectories we created, we sampled

> Fig. 4. Children show higher uniqueness in their search than adults. a) Children choose a higher fraction of unique shapes in their game (Children: mean \pm ste = 0.40 \pm 0.02, 95% CI = [0.37,0.44], Adults: mean \pm ste = 0.27 \pm 0.02, 95% CI = [0.24,0.31]. Mann-Whitney test: U(14,746) = 4737, $p < 10^{-5}$, Rank biserial correlation = 0.358). Inset, the differences in the CDF of children's vs. adults' distributions of the fraction of unique shapes. Error bars are 95% CI of each quantile difference, b) Children exploit less common categories compared to adults (Children: mean \pm ste = $0.17~\pm~0.01,~95\%$ CI =~ [0.14,0.20], Adults: mean \pm ste = 0.24 \pm 0.01, 95% CI = [0.21,0.27]. Mann-Whitney test: U(14,746) = 9307, p < 0.0004, Rank biserial correlation =



b

-0.262). Inset, the differences in the CDF of the children's vs. adults' distributions of the % of common categories. Error bars are 95% CI of each quantile difference,



Fig. 5. As children grow, they navigate in spaces that more closely resemble the adults' search space. a) Older children (age > 6 yo) save a smaller fraction of unique shapes than younger children (age \leq 6 yo). Younger children: mean \pm ste = 0.44 ± 0.02 , 95% CI = [0.40, 0.49], Older children: mean \pm ste = 0.35 \pm 0.03, 95% CI = [0.30, 0.40]. Mann-Whitney test: U(5265) = 3287, p = 0.010, Rank biserial correlation = -0.249. Inset, the differences in the CDF of children's and adults' distributions of the fraction of unique shapes. Error bars are 95% CI of each quantile difference, b) Older children save more shapes that are shared with adults than younger children (Younger children: mean \pm ste = 38.4% \pm 0.1%, 95% CI = [35.8%, 41.8%], Older children: mean \pm ste = 48.2% \pm 0.1%, 95% CI = [48.0%, 52.8%]. Mann-Whitney test: $U(10^6) =$

 10^6 , $p < 10^{-5}$, Rank biserial correlation = 0.99). Inset, the differences in the CDF of younger and older children's bootstrapped distributions of the percentage of shared shapes. Error bars are 95% CI of each quantile difference,

from the children's distribution of number of moves between chosen shapes and their timings. The thresholding algorithm then recognizes exploration and exploitation phases in the random searches, as in the children's searches, but the shapes in this process do not necessarily bear a specific meaning or a collective theme.

Indeed, when we compare the adults' distribution to the children's distribution, both show a similar scaling law, calculated as the $% \left({{\left[{{{\rm{ch}}} \right]}_{\rm{ch}}} \right)$

maximum likelihood exponent of the power-law distribution, $\alpha = 1 +$

 $n\left(\sum_{i} Log\left(\frac{x_i}{x_{min}}\right)\right)^{-1}$ (Newman, 2005) (Adults: $\alpha = 2.6 \pm 0.03, 95\%$ CI =

[2.54, 2.67], Children: $\alpha = 2.51 \pm 0.04$, 95% CI = [2.43, 2.59], Fig. 6). These two scaling exponents are very different from that of random exploration (Random exploration: $\alpha = 10.8 \pm 0.7$, 95% CI = [9.5, 12.4], Mann-Whitney test for random versus adults or children, *p*'s < 10⁻⁵, Fig. 6). In addition, while children's search behavior created shared common themes in children's categories (number of clusters in the giant component = 214), the random exploration simulations did not create any set of shared themes (zero clusters in the giant component). These findings suggest that children's exploration behavior is different from random exploration and is directed towards common shapes, like adults' search, though those shapes are different from the ones that adults



Fig. 6. Children's creative search is different from random exploration and is directed towards common shapes. The distribution of the number of times each shape was created – Adults (green), Children (orange), Random exploration (red). Adults and Children scaling law is similar and different from random exploration behavior (Adults: $\alpha = 2.6 \pm 0.03$, 95% CI = [2.54, 2.67], Children: $\alpha = 2.51 \pm 0.04$, 95% CI = [2.43, 2.59], Random exploration: $\alpha = 10.8 \pm 0.7$, 95% CI = [9.5, 12.4]). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

produce.

3. Discussion

In this study we used a novel paradigm, the "Creative Foraging Game", to map and quantify children's creative search, compare children's creative search to the creative search of adults, and to track the development of creative search through childhood.

We found marked differences in how children and adults explored the space of possibilities in the creative exploration task – differences manifested in their search strategy, the space they navigate in, and the products of their search. In particular, children spent more time exploring the space of possibilities than did adults, and when they exploited a related set of shapes, children did so less efficiently than adults. Children chose different paths to search than adults did, and the shapes they chose constitute a distinct and smaller part of the possible space of shapes. However, children also created a higher fraction of unique shapes and exploited common themes less in their search.

In addition to the differences between children and adults, we also found that creative search changes through childhood. As children develop, the space of possibilities through which they navigate grows closer to the space of adults' search. Our results suggest that this developmental arc is evident as early as 4 to 8 years of age.

Our findings also suggest that children spend more time in exploration, but their exploration is directed towards a different and smaller part of the large space of possible shapes. Most paradigms that have examined the exploration in children have done so in a constrained search space, involving choices within a relatively small space of alternatives; this precludes the possibility of finding differences in children's and adults' search space. In contrast, the CFG paradigm has a large search space, including all 36,644 possible shapes that can be produced in the grid—and in this large search space, children and adults explore different possibilities.

The results suggest that there are two different aspects to children's search compared to adults. Overall, the space of shapes children consider is smaller than the space of shapes that adults consider. This makes sense, after all – adults' common themes include letters, numbers, and other images that they have learned over many years, know well and are familiar with. However, within children's smaller search space, children appear to be more exploratory than adults. In particular, children take relatively longer paths and create more alternative shapes before they decide whether to save those shapes, and more of the shapes they do save are unique. In other words, children spend more time manipulating the squares, and the shapes they save appear to be more likely to emerge from that exploration, rather than to be conceptualized as a goal beforehand. Adults, in contrast, appear to choose a shape from

their existing mental repertoire, and then use the squares in the grid to implement that shape efficiently.

In Bayesian terms, children appear to employ a unique and narrow hyperprior on "interesting shapes" that constrains the space they choose to navigate in, while within this limited search space they consider many shapes to be similarly interesting (i.e., use broad priors, as they do in their learning of causal and social rules (Gopnik et al., 2015; Lucas et al., 2014; Seiver, Gopnik, & Goodman, 2013; Wente et al., 2019)). In this respect, children's search strategies may suggest that they resolve the exploration/exploitation trade-off differently than adults – while adults cover a larger region of the search space and restrict their exploration within it, children cover a smaller region of the search space but increase their exploration within it.

An alternative explanation for children's higher exploration, the less efficient search they exhibit, and their higher fraction of unique shapes could be that they are simply searching randomly. However, simulating random agents from children's distributions of trajectories and timings suggests that children's search is not random but rather directed at specific shapes, though different shapes than adults. This finding aligns with recent findings in children's reward exploration, which indicated similar degrees of randomness (inverse temperature parameter) in children's and adults' behavior (Schulz et al., 2019). Another possible explanation for the differences in children's search could be that children are generally less efficient in their search. Yet, if this was true, the ratio of exploration efficiency to exploitation efficiency would have been similar between children and adults, which is not the case. Children are less efficient in their exploitation, and not just less efficient in general.

Previous research used reinforcement learning paradigms to study both adults' and children's search for an external reward (Decker et al., 2016; Nussenbaum & Hartley, 2019; Schulz et al., 2019; Wu et al., 2018), where the objective function is well-defined, externally-motivated, and with well-understood algorithms for (approximately) solving the search problem. In contrast, the CFG paradigm presents a loosely defined, internally-motivated search with no clear algorithms to solve it. While the two paradigms study different search problems and the exploration behavior they support is different in nature, both lines of study suggest that exploration is not based on a purely random decision making but rather is directed towards either gain in information in the reinforcement learning paradigms or specific regions of the search space in the CFG. Further study of the two paradigms together can reveal more of the developing nature of human exploration in these different domains.

Our findings cannot explain the source for the different balance between explore and exploit that children employ in the CFG paradigm. It may be that children and adults interpret the task differently - while adults impose a higher-level goal on the task (e.g., find as many interesting shapes as possible), children might interpret it as play – a free, unconstrained exploration. Changing the context or goal of the task might make adults look more like children or vice-versa. In addition, our findings cannot discriminate between the algorithmic and representational aspects of children's search. Does their higher exploratory behavior stem from a difference in the algorithm children use as they search, from the way they conceptually organize the space of shapes, or both? Further experiments and computational work are needed to disentangle these two possibilities.

Despite these limitations, the CFG task allows a fine-grained mapping of children's creative search strategies, the space they choose to navigate in, and their products of search. In all these aspects, children show marked differences from adults' search. These differences highlight the developmental arc of the unique human ability to be creative.

CRediT authorship contribution statement

Yuval Hart: Conceptualization, Methodology, Formal analysis, Visualization, Writing – original draft, Writing – review & editing. Eliza Kosoy: Conceptualization, Methodology, Investigation, Writing – original draft, Writing – review & editing. **Emily G. Liquin:** Conceptualization, Methodology, Investigation, Writing – original draft, Writing – review & editing. **Julia A. Leonard:** Methodology, Investigation, Writing – original draft, Writing – review & editing. **Allyson P. Mackey:** Conceptualization, Methodology, Writing – original draft, Writing – review & editing, Supervision. **Alison Gopnik:** Conceptualization, Methodology, Writing – original draft, Writing – review & editing, Supervision. **Alison Gopnik:** Conceptualization, Methodology, Writing – original draft, Writing – review & editing, Supervision.

Acknowledgements

YH thanks the ISF (grant no. 3081/21). AG thanks DARPA (grant no. 047498-002), John Templeton Foundation (grant no. 61475), and Templeton World Charity Foundation (grant no. TWCF0434). AM Thanks the Jacobs Foundation Research Fellowship. EGL thanks NSF Graduate Research Fellowship Program.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.cognition.2022.105102.

References

- Amabile, T. M. (1985). Motivation and creativity: Effects of motivational orientation on creative writers. *Journal of Personality and Social Psychology*, 48(2), 393–399. https:// doi.org/10.1037/0022-3514.48.2.393
- Auer, P. (2002). Using confidence bounds for exploitation-exploration trade-offs. Journal of Machine Learning Research, 3(Nov), 397–422.
- Blanco, N. J., & Sloutsky, V. M. (2020). Systematic exploration and uncertainty dominate young children's choices. *Developmental Science*., Article e13026. https://doi.org/ 10.1111/desc.13026
- Bonawitz, E., Denison, S., Griffiths, T. L., & Gopnik, A. (2014). Probabilistic models, learning algorithms, and response variability: Sampling in cognitive development. *Trends in Cognitive Sciences*, 18(10), 497–500. https://doi.org/10.1016/j. tics.2014.06.006
- Bonawitz, E. B., van Schijndel, T. J. P., Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive Psychology*, 64(4), 215–234. https://doi.org/10.1016/j.cogpsych.2011.12.002
- Christakou, A., Gershman, S. J., Niv, Y., Simmons, A., Brammer, M., & Rubia, K. (2013). Neural and psychological maturation of decision-making in adolescence and young adulthood. *Journal of Cognitive Neuroscience*, 25(11), 1807–1823. https://doi.org/ 10.1162/jocn a 00447
- Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical transactions of the Royal Society of London Series B, Biological Sciences*, 362(1481), 933–942. https://doi.org/10.1098/rstb.2007.2098
- Cook, C., Goodman, N. D., & Schulz, L. E. (2011). Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition*, 120(3), 341–349. https:// doi.org/10.1016/j.cognition.2011.03.003
- Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095), 876–879. https://doi.org/10.1038/nature04766
- Decker, J. H., Otto, A. R., Daw, N. D., & Hartley, C. A. (2016). From creatures of habit to goal-directed learners. *Psychological Science*, 27(6), 848–858. https://doi.org/ 10.1177/0956797616639301
- Frank, M. J., Doll, B. B., Oas-Terpstra, J., & Moreno, F. (2009). Prefrontal and striatal dopaminergic genes predict individual differences in exploration and exploitation. *Nature Neuroscience*, 12(8), 1062–1068. https://doi.org/10.1038/nn.2342
- Gershman, S. J. (2018). Deconstructing the human algorithms for exploration. Cognition, 173, 34–42. https://doi.org/10.1016/j.cognition.2017.12.014
- Girvan, M., & Newman, M. E. J. (2002). Community structure in social and biological networks. Proceedings of the National Academy of Sciences, 99(12), 7821–7826. https://doi.org/10.1073/pnas.122653799
- Gopnik, A. (2020). Childhood as a solution to explore–exploit tensions. Philosophical Transactions of the Royal Society B: Biological Sciences, 375(1803), 20190502. https:// doi.org/10.1098/rstb.2019.0502
- Gopnik, A., Griffiths, T. L., & Lucas, C. G. (2015). When younger learners can be better (or at least more open-minded) than older ones. *Current Directions in Psychological Science*, 24(2), 87–92. https://doi.org/10.1177/0963721414556653
- Gopnik, A., O'Grady, S., Lucas, C. G., Griffiths, T. L., Wente, A., Bridgers, S., ... Dahl, R. E. (2017). Changes in cognitive flexibility and hypothesis search across human life history from childhood to adolescence to adulthood. *Proceedings of the National Academy of Sciences*, 114(30), 7892–7899. https://doi.org/10.1073/ pnas.1700811114
- Hart, Y., Goldberg, H., Striem-Amit, E., Mayo, A. E., Noy, L., & Alon, U. (2018). Creative exploration as a scale-invariant search on a meaning landscape. *Nature Communications*, 9(1), 5411. https://doi.org/10.1038/s41467-018-07715-8

Y. Hart et al.

Hart, Y., Mayo, A. E., Mayo, R., Rozenkrantz, L., Tendler, A., Alon, U., & Noy, L. (2017). Creative foraging: An experimental paradigm for studying exploration and discovery. *PLoS One*, *12*(8), Article e0182133. https://doi.org/10.1371/journal. pone.0182133

- Hauser, T. U., Iannaccone, R., Walitza, S., Brandeis, D., & Brem, S. (2015). Cognitive flexibility in adolescence: Neural and behavioral mechanisms of reward prediction error processing in adaptive decision making during development. *NeuroImage*, 104, 347–354. https://doi.org/10.1016/j.neuroimage.2014.09.018
- Hills, T. T., Todd, P. M., & Goldstone, R. L. (2008). Search in external and internal spaces. Psychological Science, 19(8), 802–808. https://doi.org/10.1111/j.1467-9280.2008.02160.x
- Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., & Couzin, I. D. (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*, 19(1), 46–54. https://doi.org/10.1016/j.tics.2014.10.004
- Kirkpatrick, S., Gelatt, C. D., & Vecchi, M. P. (1983). Optimization by simulated annealing. Science, 220(4598), 671–680. https://doi.org/10.1126/ science.220.4598.671
- Legare, C. H. (2012). Exploring explanation: Explaining inconsistent evidence informs exploratory, hypothesis-testing behavior in young children. *Child Development*, 83(1), 173–185. https://doi.org/10.1111/j.1467-8624.2011.01691.x
- Liquin, E., & Gopnik, A. (2020). Children are more exploratory and learn more than adults in an approach-avoid task. PsyArXiv. https://doi.org/10.31234/osf.io/6watc
- Lucas, C. G., Bridgers, S., Griffiths, T. L., & Gopnik, A. (2014). When children are better (or at least more open-minded) learners than adults: Developmental differences in learning the forms of causal relationships. *Cognition*, 131(2), 284–299. https://doi. org/10.1016/j.cognition.2013.12.010
- Lucas, C. G., & Griffiths, T. L. (2010). Learning the form of causal relationships using hierarchical Bayesian models. *Cognitive Science*, 34(1), 113–147. https://doi.org/ 10.1111/j.1551-6709.2009.01058.x
- Mata, R., Wilke, A., & Czienskowski, U. (2013). Foraging across the life span: Is there a reduction in exploration with aging? Frontiers in Neuroscience, 7. https://doi.org/ 10.3389/fnins.2013.00053

Newman, M. E. J. (2005). Power laws, Pareto distributions and Zipf's law. Contemporary Physics, 46(5), 323–351. https://doi.org/10.1080/00107510500052444

- Nussenbaum, K., & Hartley, C. A. (2019). Reinforcement learning across development: What insights can we draw from a decade of research? *Developmental Cognitive Neuroscience*, 40, Article 100733. https://doi.org/10.1016/j.dcn.2019.100733
- Oh, Y., Chesebrough, C., Erickson, B., Zhang, F., & Kounios, J. (2020). An insight-related neural reward signal. *NeuroImage*, 214, Article 116757. https://doi.org/10.1016/j. neuroimage.2020.116757
- Palminteri, S., Kilford, E. J., Coricelli, G., & Blakemore, S.-J. (2016). The computational development of reinforcement learning during adolescence. *PLoS Computational Biology*, 12(6), Article e1004953. https://doi.org/10.1371/journal.pcbi.1004953
- Pyke, G. H. (1984). Optimal foraging theory: A critical review. Annual Review of Ecology and Systematics, 15, 523–575.
- Rousselet, G. A., Foxe, J. J., & Bolam, J. P. (2016). A few simple steps to improve the description of group results in neuroscience. *European Journal of Neuroscience*, 44(9), 2647–2651. https://doi.org/10.1111/ejn.13400

- Rousselet, G. A., Pernet, C. R., & Wilcox, R. R. (2017). Beyond differences in means: Robust graphical methods to compare two groups in neuroscience. *European Journal* of Neuroscience, 46(2), 1738–1748. https://doi.org/10.1111/ejn.13610
- Rovee, C. K., & Rovee, D. T. (1969). Conjugate reinforcement of infant exploratory behavior. Journal of Experimental Child Psychology, 8(1), 33–39. https://doi.org/ 10.1016/0022-0965(69)90025-3
- Ruggeri, A., Swaboda, N., Sim, Z. L., & Gopnik, A. (2019). Shake it baby, but only when needed: Preschoolers adapt their exploratory strategies to the information structure of the task. *Cognition*, 193, Article 104013. https://doi.org/10.1016/j. cognition.2019.104013
- Schulz, E., Wu, C. M., Ruggeri, A., & Meder, B. (2019). Searching for rewards like a child means less generalization and more directed exploration. *Psychological Science*, 30 (11), 1561–1572. https://doi.org/10.1177/0956797619863663
- Schulz, L. (2012). The origins of inquiry: Inductive inference and exploration in early childhood. Trends in Cognitive Sciences, 16(7), 382–389. https://doi.org/10.1016/j. tics.2012.06.004
- Schulz, L. E., & Bonawitz, E. B. (2007). Serious fun: Preschoolers engage in more exploratory play when evidence is confounded. *Developmental Psychology*, 43(4), 1045–1050. https://doi.org/10.1037/0012-1649.43.4.1045
- Schulz, L. E., Gopnik, A., & Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Developmental Science*, 10(3), 322–332. https://doi.org/10.1111/j.1467-7687.2007.00587.x
- Seiver, E., Gopnik, A., & Goodman, N. D. (2013). Did she jump because she was the big sister or because the trampoline was safe? Causal inference and the development of social attribution. *Child Development*, 84(2), 443–454. https://doi.org/10.1111/ i.1467-8624.2012.01865.x
- Somerville, L. H., Sasse, S. F., Garrad, M. C., Drysdale, A. T., Abi Akar, N., Insel, C., & Wilson, R. C. (2017). Charting the expansion of strategic exploratory behavior during adolescence. *Journal of Experimental Psychology. General*, 146(2), 155–164. https://doi.org/10.1037/xge0000250
- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science (New York, N.Y.), 348*(6230), 91–94. https://doi.org/ 10.1126/science.aaa3799
- Sumner, E., Li, A. X., Perfors, A., Hayes, B., Navarro, D., & Sarnecka, B. W. (2019). The exploration advantage: Children's instinct to explore allows them to find information that adults miss. [Preprint]. *PsyArXiv.* https://doi.org/10.31234/osf.io/h437v
- Sutton, R. S., & Barto, A. G. (2018). Reinforcement learning: An introduction (2nd ed.). Bradford Books.
- Teodorescu, K., & Erev, I. (2014). On the decision to explore new alternatives: The coexistence of under- and over-exploration. *Journal of Behavioral Decision Making*, 27 (2), 109–123. https://doi.org/10.1002/bdm.1785
- Wente, A. O., Kimura, K., Walker, C. M., Banerjee, N., Flecha, M. F., MacDonald, B., ... Gopnik, A. (2019). Causal learning across culture and socioeconomic status. *Child Development*, 90(3), 859–875. https://doi.org/10.1111/cdev.12943
- Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2018). Generalization guides human exploration in vast decision spaces. *Nature Human Behaviour*, 2(12), 915–924. https://doi.org/10.1038/s41562-018-0467-4